Teraflop Visualization

Arthurine R. Breckenridge
Computer Architectures Department
Sandia National Laboratories, New Mexico, US
arbreck@sandia.gov

Abstract

The key to insight is coupling the power of the computer with unique skills of the human. At Sandia National Laboratories' Interaction Laboratory, we call this teraflop visualization. We are concentrating research in three main areas: 1) using the computer as a facility for authoring content, 2) adding the physics to model real behaviors, and 3) allowing the human to utilize the improved precision and resolution provided by this new class of compute power.

Keywords

Scientific Visualization, Virtual Environments, Parallel Algorithms, Haptics, Teraflop, 3D widgets, human computer interface

Introduction

> *"The primary research instrument of the sciences of complexity is the computer. Ever since the rise of modern science three centuries ago, the instruments of investigation such as telescopes and microscopes were analytic and promoted the reductionalist views of science. Physics, because it dealt with the smallest and most reduced entities, was the most fundamental science. From the laws of physics, one can deduce the laws of chemistry, then of life, and so on up the ladder. The computer, with its ability to manage enormous amounts of data and simulate reality, provides a new window on that view of physics. We may begin to see reality differently simply because the computer produces knowledge and insight differently from the traditional analytic instruments. It provides a different angle on reality."*
>
> *Heinz Pagels, The Dreams of Reason, 1988.*

Ten years later, Sandia National Laboratories researchers are providing this different angle on reality by addressing two important concepts: a) how do we get knowledge from the computer into the human, and b) how do we get insight from the human into the computer? The answers inevitably leads one to computer graphics, because the eye is the human's broad-bandwidth information channel. This is now called "scientific visualization." It has only been recently, after video games and educational software were established industries, that the scientific world has accepted the idea that computer graphics can help them understand their vast quantities of numbers. At Sandia's Interaction Laboratory, we do not want another decade to pass before our scientists see what was obvious to every kid the first time he/she touched a video game **--** the power of interactive graphics. To go one step further, we want to use all the channels of communicating with the human being that the mind already knows how to interpret. Ivan Sutherland (often called the father of computer graphics) stated you create the mathematical model of a virtual world in the computer then you want this model to look, feel, and sound as much as possible like a real world to the human mind coupled with the it. [Sutherland, 1965] This is the definition we use for “virtual environments.” (VE s) Taking this definition one step further is teraflop visualization where the computing model and the virtual environment are an integral whole needed to provide computing for insight.

In order to implement teraflop visualization we are concentrating research in three main areas: 1) authoring content (the computer) , 2) adding behaviors (the physics) , and 3) enhancing precision and resolution (the human).

1 Authoring Content

The first area of research focus is using the computer as a facility for authoring content. Having participated in decades of research, Sandia researchers are convinced that today’s computers empower us to build sophisticated models of complex natural phenomena and to explore them from new perspectives. Noting that the greatest scientific revolutions of Newton’s times were empowered by the new mathematical capabilities of the calculus and other analytic equations. We believe that the power of modern computers to represent complex mathematical models can leverage yet further exponential progress in science.

In Sandia’s Interaction Laboratory, the success stories of insight gained by using VEs are numerous. Our simulation codes are working on problems that are too hard to do by machine algorithms alone and require human insight, and are too hard to do by human insight alone and require a lot of calculations. In other words, a terabyte of information has become easy to generate; indeed, datasheets of this size are regularly being produced on the massively parallel Intel teraflop supercomputer located at Sandia. More recently, clusters of personal computers have begun to reach the gigaflop scale of computation. How can a scientist hope to comprehend this volume of data effectively?

Effective ways to answer this question are:

- Building parallel virtual environment hardware

  The Intel Teraflop machine provides an unique opportunity to author content. The visualization shown in Figure 1 is one of the result files calculated with 4096 processors during the initial benchmark of the machine. Since the Teraflop is, simplistically, a collection of personnel computer processors, one could, if able to overcome memory, disk and other bottlenecks, estimate that it would have take almost 10 years to complete this same problem on a single PC [i.e., (4096 processors X 18 hours) / (24 hours in a day X 365 days in a year)].

Virtual environments consume CPU cycles. Or, VE eat MIPS for breakfast. All who have ventured into the realm of real-time multi-dimensional simulations have discovered that VE requires a hybrid of general purpose and special processors to create even a crude approximation of reality. However, the cost of computer chips keep dropping as their power increases, which is making it possible to link hundreds, even thousands, of computer chips in special systems. Today, computing speeds are measured in MIPS (millions of instructions per second) and strange new rates such as “teraflops”.

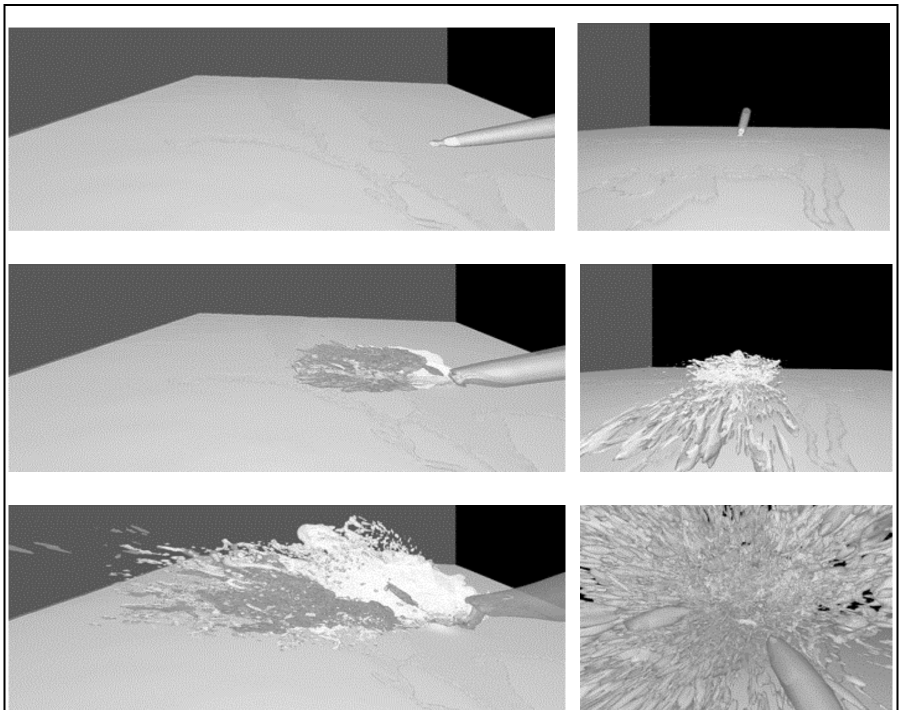

Figure 1. The 100 Million structured cell data set of a comet hitting off the coast of New York City at 3 time intervals.

The solution is to develop a general pathway for VE that can utilize the extremely high rates of processing speed in the computer architecture known as “massive parallelism,” in which many instead of one is devoted to the same task, which is logically parceled out in some manner. All machines will be based on commodity off the shelf products. (The EIGEN/VR project which is exploring this concept is named CPLANT.[1])

- Building integrated simulation and virtual environment software

  The Intel Teraflop machine provides an unique opportunity to implement the VE foresight provided by Fredrick Brooks (often called the father of virtual reality), which he proposed even before computing hardware was powerful enough to achieve his vision. Brooks believes the use of computer systems for intelligent amplification (IA) is much more powerful today, and will continue to be at any given point, than the use of the computers for artificial intelligence (AI) where the objective is to replace the human mind by the machine and its simulation codes. The virtual environment IA objective is to build systems that amplify the human mind by providing it with computer-based auxiliaries that do the things that the mind has trouble doing. Brooks offers several areas of in which human minds are more powerful than any computer algorithm designed. The computer's strength is in storing massive amounts of data and recalling the data without forgetting it. The human strength is to recall data at the appropriate time, in a reference to a completely different subject, that we suddenly see to be meaningful. Human evaluation and recall are very hard processes to translate into a well-defined computer data retrieval command. Humans are more skilled than computers at pattern recognition, whether visual or aural. [Brooks, 1977]

  The solution is to develop a general pathway for VE that utilizes the strengths of both the human and the computer. Qualitative algorithms are needed to get a high level view of what the large data sets looks like. Quantitative algorithms are needed for the human to probe beyond sight into the raw data. Also, amplifying algorithms are being developed so the computer can focus on preset reduction criteria while the human has a mechanism to steer the computation based on his/her evaluation and pattern recognition strengths. (The EIGEN/VR project which is exploring this concept is named Parallel Algorithms.[1])

2 Adding Behaviors

The second area of research focus at Sandia is adding the physics to model real behaviors. Simply specifying an object "static" 3D-geometry is not sufficient for analysis. The object's behavior in the VE involves changes in position, collisions, grasping, scaling, surface deformations, etc. Virtual objects also need to be modeled physically by specifying their mass, weight, surface texture (smooth or rough), deformation mode, etc. These features are merged with geometrical modeling and behavior laws to form a more realistic VE. As an example Figure 2 shows the crating results from the high velocity impact physics of the comet.

In the Interaction Laboratory, we not only concentrate on the behaviors coded into the simulation, we also focus on the human behaviors needed to interact with the data and the computer. There is more to modeling the real world than drawing a nice image. We make all of our scientists handicapped when interacting with a computer if we only give them their sense of sight and no tools to manipulate the images. How can a scientist examine the volume of data and its behaviors?

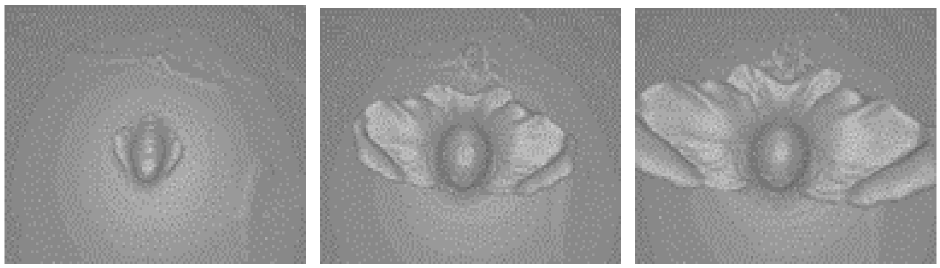

Figure 2. Illustrates the impact of the comet on the ocean floor.

Effective ways to answer this question are:

- Mapping force and tactile (haptics) feedback

Our eyes may be the main broad-bandwidth information channel, but we quickly realize the importance of our other senses when we operate without them in the VE. Much progress has been made in adding visual and auditory feedback possibly since the human receptors are localized. This is not true for haptic feedback. Not only are the receptors not localized, they are not consistent over the body surface. A great impact can be made by starting with the fingers. The maximum tactile sensor density is at the fingertips, where the density is up to 135 sensors/cm^2 [Kokjer, 1987].

The sensation associated with haptic sensors also involves the melange of senses we lump together under the category of "touch." Haptic is not strictly tactile in the same way one's fingertips convey tactile information about the outside world. Haptic tasks also include tasks that use the body's sense of proprioception that informs us about the position of our own limbs in relation to one another and to the space around us. [Boff, 1986] We hardly notice the mind's fast, silent information-processing and fine muscular coordination skill that enables you to move your hand in exactly the right direction when you decide to reach for a glass of water. A ballet dancer is a virtuoso of proprioception. Haptics involves both proprioceptive and tactile senses, in concert with other senses.

The solution is to add as many natural behaviors as possible to the VE. The VE is enhanced by adding auditory cues and haptic feedback. Sound but especially touch are critical elements in developing effective human/computer interaction. (The EIGEN/VR project which is exploring this concept is named FLIGHT.[1])

- Mapping human imperceptible phenomena.

  We have certain built-in senses such as vision, hearing, and touch, but there are many phenomena which are completely imperceptible to us. Some examples are x-rays, radioactivity, and electricity. One could say we lack the senses to perceive these things. But, by using electronic sensors and computer interfaces, we can make these imperceptible phenomena visible, or audible or touchable. One example would be magnetic fields which are traditionally visualized with vector arrows in a 2D visualization (called the hedgehog). In 3D, the arrows are very confusing due to occlusion problems. A human can see iron filings moving or the effect of a magnetic field but can not see the magnetic force itself. We must look beyond vision to convey these imperceptible phenomena.

  The solution is to use develop innovative interfaces that map very hard physics into the appropriate sense, which will enhance comprehension. (The EIGEN/VR project which is exploring this concept is named HEDGEHOG.[1])

3 Utilizing Precision and Resolution

The third area of research focus at Sandia concerns allowing the human to utilize the improved precision and resolution. The Teraflop machine provides an unique opportunity to increase the resolution and precision of the simulation. Figures 3a and 3b illustrate the difference in detail of a 54 Million cell volume representation of the comet computed in 1996 and a 100 M cell surface representation from 1997. Besides the well known limitations such as system latency and display resolution, the precise manipulation of virtual objects is very difficult. Why add resolution and precision if it must be removed in order to visualize the data?

Effective ways to answer this question are:

- Upgrading the settings

  Thirty thousand years ago, many of the ochre paintings on limestone cave walls at Lascaux were painted in a precisely distorted manner in order to give the rendering a 3D appearance. The art itself is only part of the experience, often rather a small part at that. The ochre figures viewed out of context and in two dimensions, reproduced in books are far less interesting and informative. The setting was a major factor in the effect of the figures. Considering the technologies at the time, the people seemed to have made use of every trick, using special effect upon special effect in an effort to create a specific state of consciousness to focus the human mind's attention on the information. [Pfeiffer, 1982]

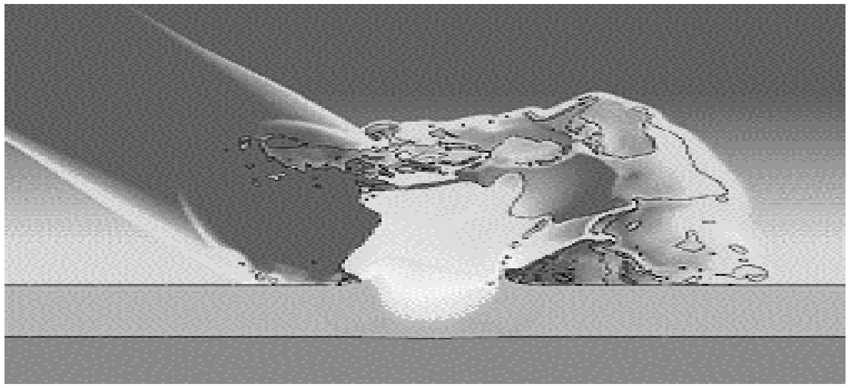

Figure 3a: 54 Million cell representation of two variables (water and comet)

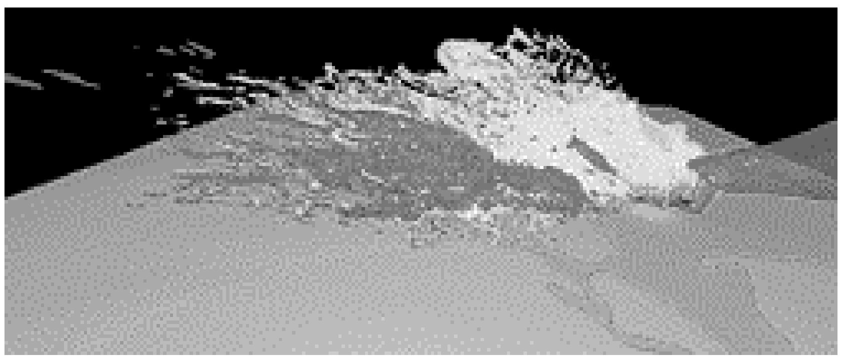

Figure 3b: 100 Million cell representation with five variables (air, water, land, temperature and comet)

Like the caves at Lascaux, recent alternative output devices such as CAVE (TM)-like environments or stereoscopic wall projections and more precise input devices for enhancing the experience [Cruz-Neira, 1992]. These environments have gone far in providing the settings for making the computer become transparent to allow focusing on the information. The next step is to reduce the cost of these tools, extend the domain space to more collaboration than a few individuals, and improve the resolution and precision of interaction. In these CAVEs, as in the caves at Lascaux, people in altered state of consciousness can be educated in the use of tools for creating and maintaining knowledge.

The solution is to provide the “setting” for focused attention whether in a CAVE, VE on the desktop, or the next generation setting. Real environments require people to be in the same space at the same time. VEs will hopefully get rid of the real space requirement allowing humans in distant settings to collaborate. We are going to use this new communication medium to take us beyond the boundaries of our minds and push our society to create and transmit information from generation to generation. (The EIGEN/VR project which is exploring this concept is named KIVA [1].)

- Upgrading from a 2D to 3D interface

Thirty thousand year ago, well not really, but about three decades have passed since devoting a single computer to the use of one person was a revolutionary idea that converged with another revolutionary idea -- the notion that the human interfaces for computers should accommodate human needs and abilities, rather than shaping human behavior to fit the demands of the computer technology. Today, the word computer brings to mind monitor, keyboard, and mouse attached to a desktop machine. VE s are supporting technologies that will explode that image.

Let's consider the current human/computer interface, or the lack thereof. To solve problems with the aid of a computer, a user must be able to easily and accurately communicate to the machine and control what need to be done. Designing the means to do so (i.e., the human/computer interface) is by no means trivial. The way in which a person enters commands or data into a computer is the input part of the interface, and the way the computer shows the user the results of the computations is the output part of the interface. In the 1950s, the first input and output devices were designed with the limitations of computers rather than the capabilities of humans in mind. A computer that accepts input only in the form of punched cards and spits out answers on a roll of paper, is an example of a user interface designed to meet the needs of computing machines.

The mouse pointing device, now ubiquitous with human/computer interfaces of all kinds, was invented in the 1960's, although it was not commercially available until the 1980s. The use of a mouse as a pointing device signalled a concrete break-through in the human-computer interface, one that moved directly to the core of VE: gestural input as a command language. However, the mouse implementation (WIMP: Windows, Icons, Menus, and Pointers) was in a 2D environment that reduced the look and feel so any object/widget could never be more than flat object and, hence, not intuitive. Although standardization of user interfaces (Motif) has been an important development in this decade, Sandia researchers feel this is still one of the largest areas for improvement. The invention of the keyboard and mouse is better; however we can still improve on this.

The solution is stop mapping 3D functions into 2D space when at all possible. Our simulations are moving to 3D as quickly as possible; so, we should move to 3D interfaces as quickly as possible. We need a next-generation mouse, but more importantly we need "de-facto standard" 3D widget software that will operate with the peripherals of the future based on 6 DOF (x, y, z, row, pitch, and yaw) plus force, plus sound, etc. Today, our scientists are not enamored with the encumbering devices like "goggles and gloves". But as the devices become more natural and the interfaces more comfortable, familiar, and effortless (like the telephone), we can unconsciously subtract our awareness of the VE peripherals, and focus on the function -- the information. (The EIGEN/VR project which is exploring this concept is named VRML widgets[1])

Changing the angle of reality

In summary, virtual environments offer the potential for another technological revolution like the telephone and television. Sandia researchers think VE will be the technological medium for advancing computing for insight. Alexander Graham Bell thought the telephone would be useful as a way to pipe music to people. The developers of radio and television thought that their devices would launch a world of 2-way communications to replace the telephone. Ideas abounded about possible implications of VE s. In retrospect, we now understand something about the way telephones and television expanded far beyond the expectations of their inventors. In fact, it is interesting to compare the introduction of television with the introduction of virtual environments. Here was a technology fascinating to view, though limited in content and image quality. Crowds formed wherever it was initially demonstrated. In wide spread use, the main application of television turned out to be the introduction of the concept of telepresence, the sense of being there through the eyes of a camera in realtime. In wide spread use, the prime application for coupling computers with VE s is still to be determined.

It is hoped that reading about the focus areas of Sandia's Interaction Laboratory staff has been useful and informative. The VE peripherals will change (i.e., holovideos and autostereosopic displays) and our environment will become more natural. The author has no doubt that teraflop visualization (coupling computers and VE s) will influence our society greatly in the not too distant future. As John Thomas (often called the father of innovative interfaces), put it:

> "*VE, at best, will not only affect the world we live in, it may help ensure we have a world to live in." [Thomas, 1993]*

Boff, K. R., L. Kaufman and J.P. Thomas, Eds. Handbook of Perception and Human Performance, New York: John Wiley and Sons, 1986.

Brooks, F.P, Jr., "The Computer Scientist as Toolsmith: Studies in Interactive Computer Graphics," Information Processing 77, B. Gilchrist, ed., Amsterdam: North-Holland, 1977, pp. 625-634.

Cruz-Neira, C., Sandin, D.J., DeFanti, T.A., Kenyon, R., and Hart, J.C. The CAVE, Audio Visual Experience Automatic Virtual Environment. Communications of the ACM, June 1992.

Kokjer, K., "The Information Capacity of the Human Fingertip," IEEE Transactions on Systems, Man and Cybernetics, 1987, Vol. SMC-17, No. 1, pp. 100-102.

Pagels, H., The Dreams of Reason, New York: Simon & Schuster, 1988.

Pfeiffer, J. E., The Creative Explosion: An Inquiry into the Origins of Art and Religion, Ithaca, NY: Cornell University Press, 1982.
Sutherland, I., “The Ultimate Display,” Proceedings of the IFIP Congress, 1965, pp. 506-508.
Thomas, J., “The Shape of Things to Come,” Proceedings of VR Systems ‘93 Conference, New York, NY, pp. 342-346, March, 1993.

[1]This keynote address at Visualization Environment ‘98 was meant to identify trends in VE. For more specific information about Sandia’s work in removing these limitations please see details on our web pages at http://www.cs.sandia.gov/ILAB. All sub-projects mentioned are part of the EIGEN/VR teraflop visualization project.

This work was supported by Sandia National Laboratories, a multi-program laboratory operated by Sandia Corporation (a Lockheed Martin company) for the United States Department of Energy under contract DE-AC04-94AL85000.